



energie atomique • énergies alternatives



## **Bruyères-le-Châtel – France**

**January 2012**

### **CEA's largest HPC file system reaches 11PB**

BRUYERES-LE-CHATEL, France, Jan. 24 2012 -- The Military Applications Department of the French Alternative Energies and Atomic Energy Commission (CEA/DAM) today announced that it doubled the capacity of its largest HPC production file system. With over 11 Petabytes of usable storage space, this Lustre v2 file system is one of the largest single namespace system ever deployed. A Petabyte (PB) is a quadrillion bytes, a thousand times more that today's current desktop computer storage space.

Starting in September 2010, the storage capacity has been gradually increased, to reach today 75% of TERA computing center's global Lustre storage space and bandwidth (200GB/s), a step forward to its final configuration of 15PB planned this year.

With 768 16TB Object Storage Targets (OST), TERA global file system consists in 70 bullx S6030 Lustre servers, a Voltaire Infiniband QDR storage fabric and 89 NetApp (formerly LSI) E5400-60 disk arrays. Tera 100, Europe's first supercomputer to break the Petaflops barrier ([http://www-hpc.cea.fr/docs/cp-Tera100-091110\\_VE1.pdf](http://www-hpc.cea.fr/docs/cp-Tera100-091110_VE1.pdf)), and other ancillary systems such as post-processing clusters, mount this file system through more than 50 Lustre routers.

### **System administration**

To optimize deployment time, availability and to simplify daily administration tasks, CEA engineers use a common set of open libraries and tools accross CEA's computing centers: TERA, TGCC and CCRT (<http://www-hpc.cea.fr>). Shine, an open-source Lustre administration tool developed at CEA, manages all Lustre components, allowing configuration of servers, routers and thousands of clients (<http://lustre-shine.sourceforge.net>).

## **Data workflow**

With such a large amount of data, traditional tools usually used to scan and maintain a file system are notoriously inefficient. CEA developed Robinhood (<http://robinhood.sourceforge.net>), an open-source policy engine designed to address this specific issue. It uses Lustre MDT changelogs (a Lustre v2 feature) to update an internal database reflecting the file system state. CEA also actively contributes to Lustre development, with features such as OST pools, and is the main developer for the forthcoming Lustre-HSM binding feature (Lustre 2.x).

CEA also develops NFS-Ganesha (<http://nfs-ganesha.sourceforge.net>), a user-space NFS server capable of using the Lustre API to export very large file system to non Lustre clients, such as smaller systems or individual workstations.

## **About the CEA**

The French Alternative Energies and Atomic Energy Commission (CEA) leads research, development and innovation in four main areas: low-carbon energy sources, global defense and security, information technologies and healthcare technologies. The CEA's leadership position in the world of research is built on a cross-disciplinary culture of engineers and researchers, ideal for creating synergy between fundamental research and technology innovation. With its 15,600 researchers and collaborators, it has internationally recognized expertise in its areas of excellence and has developed many collaborations with national and international, academic and industrial partners.

Information about HPC at CEA can be found at <http://www-hpc.cea.fr/index-en.htm>